

DEGREES OF STRETCHED KOSTKA COEFFICIENTS

TYRRELL B. MCALLISTER

ABSTRACT. Given a partition λ and a composition β , the *stretched Kostka coefficient* $\mathcal{K}_{\lambda\beta}(n)$ is the map $n \mapsto K_{n\lambda, n\beta}$ sending each positive integer n to the Kostka coefficient indexed by $n\lambda$ and $n\beta$. Kirillov and Reshetikhin [11] have shown that stretched Kostka coefficients are polynomial functions of n . King, Tollu, and Toumazet have conjectured that these polynomials always have nonnegative coefficients [7], and they have given a conjectural expression for their degrees [8].

We prove the values conjectured by King, Tollu, and Toumazet for the degrees of stretched Kostka coefficients. Our proof depends upon the polyhedral geometry of Gelfand–Tsetlin polytopes and uses tilings of GT-patterns, a combinatorial structure introduced in [3].

1. INTRODUCTION

Kostka coefficients are important numbers appearing in many branches of mathematics, including representation theory, the theory of symmetric functions, and algebraic geometry (see, *e.g.*, [4, 12, 14] and references therein). Given a dominant weight λ and a weight β of the Lie algebra $\mathfrak{gl}_r(\mathbb{C})$, the Kostka coefficient $K_{\lambda\beta}$ is the dimension of the weight subspace with weight β of the irreducible representation V_λ of $\mathfrak{gl}_r(\mathbb{C})$ [5]. In the theory of symmetric functions, Kostka coefficients are defined by the expansion of Schur functions s_λ into monomials. That is, given a partition λ of $N \in \mathbb{Z}_{\geq 0}$ with r parts, we have that

$$(1) \quad s_\lambda = \sum_{\substack{\text{compositions} \\ \beta \text{ of } N}} K_{\lambda\beta} x^\beta,$$

where $x^\beta = x_1^{\beta_1} x_2^{\beta_2} \cdots x_r^{\beta_r}$.

Since the parameters defining a Kostka coefficient are themselves vectors, they may be “stretched” by an integer scaling factor n . This procedure defines a function $n \mapsto K_{n\lambda, n\beta}$, which, following [7], we call a *stretched Kostka coefficient*. We denote this function by $\mathcal{K}_{\lambda\beta}(n) = K_{n\lambda, n\beta}$. Kirillov and Reshetikhin have shown that $\mathcal{K}_{\lambda\beta}(n)$ is a polynomial function of n [11] (see also [3, Proposition 2.6]). King, Tollu, and Toumazet have conjectured that these polynomials have only positive coefficients [7], and they have given a conjectural expression for the degree of $\mathcal{K}_{\lambda\beta}(n)$ [8]. The main result of this note (Theorem 1.1 below) is that the stretched Kostka coefficients indeed have the degrees conjectured in [8].

For our purposes, it will suffice to assume that λ is a partition and β is a composition of the same length and size as λ . That is, we take $\lambda = (\lambda_1, \dots, \lambda_r) \in \mathbb{Z}_{\geq 0}^r$

2000 *Mathematics Subject Classification.* Primary 17B10; Secondary 52B12.

Key words and phrases. Kostka coefficient; representation theory; Gelfand–Tsetlin polytope.

Research supported by NSF VIGRE Grant No. DMS-0135345 and by NWO Mathematics Cluster DIAMANT.

such that $\lambda_1 \geq \dots \geq \lambda_r$ and $\beta = (\beta_1, \dots, \beta_r) \in \mathbb{Z}_{\geq 0}^r$ such that $\sum_i \beta_i = \sum_i \lambda_i$. The Kostka coefficients $K_{\lambda\beta}$ are indexed by such pairs (λ, β) . We sometimes write $\lambda = (\kappa_1^{v_1}, \dots, \kappa_m^{v_m})$ to indicate that λ has v_p parts equal to κ_p for $1 \leq p \leq m$. The use of this notation always presumes that $v_p \geq 1$ for $1 \leq p \leq m$.

Let $|\lambda| = \sum_i \lambda_i$ denote the size of λ . Given an arbitrary sequence of nonnegative integers $\beta = (\beta_1, \dots, \beta_r) \in \mathbb{Z}_{\geq 0}^r$, let $\bar{\beta} = (\bar{\beta}_1, \dots, \bar{\beta}_r)$ be the unique partition that may be produced by permuting the terms of β . We say that λ *dominates* β , denoted $\beta \trianglelefteq \lambda$, if $|\lambda| = |\beta|$ and $\sum_{k=1}^i \lambda_k \geq \sum_{k=1}^i \bar{\beta}_k$ for $1 \leq i < r$. If, in addition, $\sum_{k=1}^i \lambda_k > \sum_{k=1}^i \bar{\beta}_k$ for $1 \leq i < r$, we write $\beta \triangleleft \lambda$, and we say that λ and β form a *primitive pair*.

Theorem 1.1 (Proved on p. 9). *Suppose that $\lambda = (\kappa_1^{v_1}, \dots, \kappa_m^{v_m}) \in \mathbb{Z}_{\geq 0}^r$ is a partition and β is a composition such that $\beta \triangleleft \lambda$. Then the degree of the stretched Kostka coefficient $\mathcal{K}_{\lambda\beta}(n)$ is given by*

$$(2) \quad \deg \mathcal{K}_{\lambda\beta}(n) = \binom{r-1}{2} - \sum_{p=1}^m \binom{v_p}{2}$$

(where we evaluate $\binom{1}{2} = 0$)¹.

Example 1.2. Let $\lambda = (4, 2, 2, 0, 0, 0) = (4^1, 2^2, 0^3)$ and $\beta = (3, 1, 1, 1, 1, 1)$. Then $\lambda \triangleleft \beta$, so Theorem 1.1 applies. In this case, we have $r = 6$, $v_1 = 1$, $v_2 = 2$, and $v_3 = 3$, so the degree of $\mathcal{K}_{\lambda\beta}(n)$ is

$$\binom{6-1}{2} - \binom{1}{2} - \binom{2}{2} - \binom{3}{2} = 6.$$

As stated, Theorem 1.1 gives the degree of a stretched Kostka coefficient only when λ and β are a primitive pair. However, Berenstein and Zelevinsky have shown that all Kostka coefficients factor into a product of Kostka coefficients indexed by primitive pairs [1]. It follows from this factorization that Theorem 1.1 suffices to describe the degrees of stretched Kostka coefficients in all cases.

The factorization of Kostka coefficients works as follows. It is well known that $K_{\lambda\beta}$ is invariant under permutations of the coordinates of β . For example, this follows from equation (1) and the fact that Schur functions are symmetric. Consequently, $\mathcal{K}_{\lambda\beta}(n) = \mathcal{K}_{\lambda\bar{\beta}}(n)$. In particular, to compute the degree of stretched Kostka coefficients, we need only consider the case where β is a partition.

Suppose that $\lambda, \beta \in \mathbb{Z}^r$ are both partitions with $|\lambda| = |\beta|$. If λ and β do not form a primitive pair, then we may write λ and β , respectively, as concatenations of partitions such that each of the partitions contained in λ forms a primitive pair with the corresponding partition contained in β . More precisely, there exists a unique sequence of integers

$$1 = i_1 < i_2 < \dots < i_s < i_{s+1} = r+1$$

such that each pair

$$\lambda^{(t)} = (\lambda_{i_t}, \dots, \lambda_{i_{t+1}-1}), \beta^{(t)} = (\beta_{i_t}, \dots, \beta_{i_{t+1}-1}), \quad 1 \leq t \leq s$$

¹Equation (2) was asserted in [9, Section 7.5]. However, the statement given there omits the condition that λ and β form a primitive pair. Without this condition, the right-hand side of (2) only gives an upper bound, as was observed in [2]. For example, take $\lambda = (4, 2, 1)$ and $\beta = (3, 3, 1)$. Then $\deg \mathcal{K}_{\lambda\beta}(n) = 0$, but the right-hand side of (2) evaluates to 1.

is primitive. We then have

$$K_{\lambda\beta} = \prod_{t=1}^s K_{\lambda^{(t)}\beta^{(t)}}.$$

This observation of Berenstein and Zelevinsky [1] is the justification for the terminology *primitive pair*. Since the set of indices i_1, \dots, i_{s+1} decomposing (λ, β) into primitive pairs does not change when we scale λ and β by a parameter n , this decomposition carries over to the stretched Kostka coefficients:

$$\mathcal{K}_{\lambda\beta}(n) = \prod_{t=1}^s \mathcal{K}_{\lambda^{(t)}\beta^{(t)}}(n).$$

Consequently, we have that

$$\deg \mathcal{K}_{\lambda\beta}(n) = \sum_{t=1}^s \deg \mathcal{K}_{\lambda^{(t)}\beta^{(t)}}(n),$$

where each term in the sum may be computed using Theorem 1.1.

This note is organized as follows. In Section 2, we discuss the polyhedral model of Kostka coefficients using Gelfand–Tsetlin polytopes. In this context, the degrees of stretched Kostka coefficients are the dimensions of these polytopes. We prove that the expression in Theorem 1.1 gives the dimension of a Gelfand–Tsetlin polytope under certain generic circumstances. In Section 3, we show that these circumstances obtain precisely when the corresponding Kostka coefficient is indexed by a primitive pair, proving Theorem 1.1.

2. DIMENSIONS OF GELFAND–TSETLIN POLYTOPES

Several combinatorial interpretations of Kostka coefficients have appeared in the literature. Most classically, $K_{\lambda\beta}$ is the number of semi-standard Young tableaux with shape λ and content β (see, *e.g.*, [13]). Of particular interest for our study is the representation of Kostka coefficients as the number of lattice points in particular families of rational polytopes. Gelfand and Tsetlin provided the first such model [6], which we describe and employ in our study below.

The theory of lattice point enumeration has proved to be a powerful tool for understanding the behavior of Kostka coefficients. For example, Billey, Guillemin, and Rassart have used vector partition functions to show that $\mathcal{K}_{\lambda\beta}(n)$ can be expressed as a *multivariate* piecewise polynomial in n and the coordinates of λ and β . They also examined factorizations of these polynomials and gave upper bounds on their degrees [2]. More recently, King, Tollu, and Toumazet have introduced K-hive polytopes [7], which they used to motivate the conjectures mentioned in the introduction. Moreover, they deduce from their model additional information about the structure of the polynomials $\mathcal{K}_{\lambda\beta}(n)$. Among other results, they provide an interpretation for the roots of $\mathcal{K}_{\lambda\beta}(n)$.

Expressing Kostka coefficients using Gelfand–Tsetlin polytopes provides a natural geometric interpretation of the polynomial $\mathcal{K}_{\lambda\beta}(n)$ and its degree. To each Kostka coefficient $K_{\lambda\beta}$, there corresponds a Gelfand–Tsetlin polytope $GT_{\lambda\beta} \subset \mathbb{R}^D$ such that $K_{\lambda\beta} = |GT_{\lambda\beta} \cap \mathbb{Z}^D|$, where $D = \binom{r+1}{2}$. As a consequence of the definition of Gelfand–Tsetlin polytopes (Definition 2.1 below), scaling λ and β by a positive integer n corresponds to dilating the polytope $GT_{\lambda\beta}$ by n :

$$\mathcal{K}_{\lambda\beta}(n) = |nGT_{\lambda\beta} \cap \mathbb{Z}^D|.$$

As mentioned in the introduction, Kirillov and Reshetikhin have shown that $\mathcal{K}_{\lambda\beta}(n)$ is in fact a polynomial function of n . As is well known from the theory of lattice point enumeration in polyhedra, if the number of lattice points in an integer dilation nP of a polytope P is a polynomial function of n , then the degree of that polynomial is the dimension of P (see, *e.g.*, [13, Theorem 4.6.25]). Our proof of Theorem 1.1 depends on this interpretation of $\deg \mathcal{K}_{\lambda\beta}(n)$ as the dimension of the Gelfand–Tsetlin polytope $GT_{\lambda\beta}$, which we now define.

Let X_r be the set of triangular arrays $(x_{ij})_{1 \leq i \leq j \leq r}$. Note that X_r inherits a normed vector space structure under the obvious isomorphism $X_r \cong \mathbb{R}^{r(r+1)/2}$. Therefore, we will be able to speak of cones, lattices, and polytopes in X_r . It is customary to depict an array in X_r by arranging its entries as follows:

$$\begin{array}{ccccccc}
 x_{1r} & & x_{2r} & & x_{3r} & \cdots & x_{rr} \\
 & \cdots & \cdots & \cdots & \cdots & \cdots & \cdots \\
 & & x_{13} & & x_{23} & & x_{33} \\
 & & x_{12} & & x_{22} & & \\
 & & & & x_{11} & &
 \end{array}$$

Within the space X_r , we define the cone of *Gelfand–Tsetlin patterns*, or *GT-patterns*, as follows. A GT-pattern is a triangular array $\mathbf{x} = (x_{ij})_{1 \leq i \leq j \leq r} \in X_r$ satisfying the *Gelfand–Tsetlin inequalities*

$$(3) \quad x_{i,j+1} \geq x_{ij} \geq x_{i+1,j+1}, \quad \text{for } 1 \leq i \leq j \leq r-1.$$

For any point $\mathbf{x} \in X_r$, we define the *highest weight* $\text{hwt}(\mathbf{x})$ of \mathbf{x} to be the top row (x_{1r}, \dots, x_{rr}) of \mathbf{x} , and we define the *weight* $\text{wt}(\mathbf{x}) = (\beta_1, \dots, \beta_r)$ of \mathbf{x} by $\beta_1 = x_{11}$ and $\beta_j = \sum_{i=1}^j x_{ij} - \sum_{i=1}^{j-1} x_{i,j-1}$ for $2 \leq j \leq r$. Thus, wt and hwt are both linear maps $X_r \rightarrow \mathbb{R}^r$.

For each $\lambda \in \mathbb{R}^r$, let GT_λ be the polytope of GT-patterns with highest weight λ :

$$GT_\lambda = \{\mathbf{x} \in X_r : \mathbf{x} \text{ is a GT-pattern and } \text{hwt}(\mathbf{x}) = \lambda\}.$$

Note that the GT inequalities (3) force the top row of a GT-pattern to be weakly decreasing, so $GT_\lambda = \emptyset$ if λ is not weakly decreasing. (The converse is also true, since if λ is weakly decreasing, then $(x_{ij})_{1 \leq i \leq j \leq r}$ with $x_{ij} = \lambda_i$ is a GT-pattern in GT_λ .) For each $\beta \in \mathbb{R}^r$, let $W_\beta \subset X_r$ be the affine subspace of points in X_r with weight β :

$$W_\beta = \{\mathbf{x} \in X_r : \text{wt}(\mathbf{x}) = \beta\}.$$

Definition 2.1. Given $\lambda, \beta \in \mathbb{R}^r$, the *Gelfand–Tsetlin polytope* $GT_{\lambda\beta}$ is the convex polytope of GT-patterns with highest weight λ and weight β :

$$GT_{\lambda\beta} = GT_\lambda \cap W_\beta.$$

Observe that each GT-pattern \mathbf{x} lives in a unique GT-polytope. We denote this polytope by $GT(\mathbf{x}) = GT_{\text{hwt}(\mathbf{x}), \text{wt}(\mathbf{x})}$.

We associate to each GT-pattern a certain combinatorial object, called a *tiling*. Tilings of the GT-patterns in a GT-polytope have a poset structure isomorphic to the face lattice of the GT-polytope. In particular, the tiling of a point \mathbf{x} provides

a straight-forward way to compute the dimension of the lowest-dimensional face of $GT(\mathbf{x})$ that contains \mathbf{x} (Theorem 2.4).

Definition 2.2. The *tiling* \mathcal{T} of \mathbf{x} is that partition of the set

$$\mathcal{I} = \{(i, j) \in \mathbb{Z}^2 : 1 \leq i \leq r, 1 \leq j \leq r\}$$

into subsets, called *tiles*, that results from grouping together those entries in \mathbf{x} that are equal and adjacent. More precisely, \mathcal{T} is that partition of \mathcal{I} such that two pairs $(i, j), (i', j')$ are in the same tile if and only if there are sequences

$$\begin{aligned} i &= i_1, i_2, \dots, i_s = i' \\ j &= j_1, j_2, \dots, j_s = j' \end{aligned}$$

such that, for each $k \in \{1, \dots, s-1\}$, we have that

$$(i_{k+1}, j_{k+1}) \in \{(i_k + 1, j_k + 1), (i_k, j_k + 1), (i_k - 1, j_k - 1), (i_k, j_k - 1)\}$$

and $x_{i_{k+1}j_{k+1}} = x_{i_kj_k}$.

In other words, the tiles are the connected components in the graph of a Gelfand–Tsetlin pattern. See Figure 2 for examples of GT-patterns and their tilings. The shading of some of the tiles in that figure is explained below.

Given a GT-pattern \mathbf{x} with tiling \mathcal{T} , we associate to \mathcal{T} (or, equivalently, to \mathbf{x}) a matrix $A_{\mathcal{T}}$ (or $A_{\mathbf{x}}$) as follows. Define the *free rows* of \mathbf{x} to be those that are neither the top nor the bottom row of \mathbf{x} . The *free tiles* T_1, \dots, T_s of \mathcal{T} are those tiles in \mathcal{T} that intersect only free rows of \mathbf{x} —i.e., those tiles that contain neither $(1, 1)$ nor (i, r) for $1 \leq i \leq r$. The remaining tiles are the *non-free tiles*. The order in which the free tiles are indexed will not matter for our purposes, but, for concreteness, we adopt the convention of indexing the free tiles in the order that they are initially encountered as the entries of \mathbf{x} are read from left to right and bottom to top. Define the *tiling matrix* $A_{\mathcal{T}} = A_{\mathbf{x}} = (a_{jk})_{1 \leq j \leq r-2, 1 \leq k \leq s}$ by

$$a_{jk} = \#\{i : (i, j+1) \in T_k\}.$$

That is, a_{jk} counts the number of entries in the j th free row of \mathbf{x} that are contained in the free tile T_k . While a different choice of order for the free tiles would result in a tiling matrix with permuted columns, this will be immaterial because we will ultimately be interested only in the rank of $A_{\mathcal{T}}$. The terminology *free tile* represents the fact that the entries in the non-free tiles of \mathbf{x} are fixed once the weight and highest weight of \mathbf{x} are fixed.

Example 2.3. Two GT-patterns and their tilings are given in Figure 2. The unshaded tiles are the free tiles. Index the free tiles in the order that they are initially encountered as the entries of \mathbf{x} are read from left to right and bottom to top. Then the associated tiling matrices are respectively

$$\begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 2 & 2 & 0 \\ 1 & 1 & 1 \end{bmatrix}.$$

The primary motivation for introducing tilings of GT-patterns is the following Theorem:

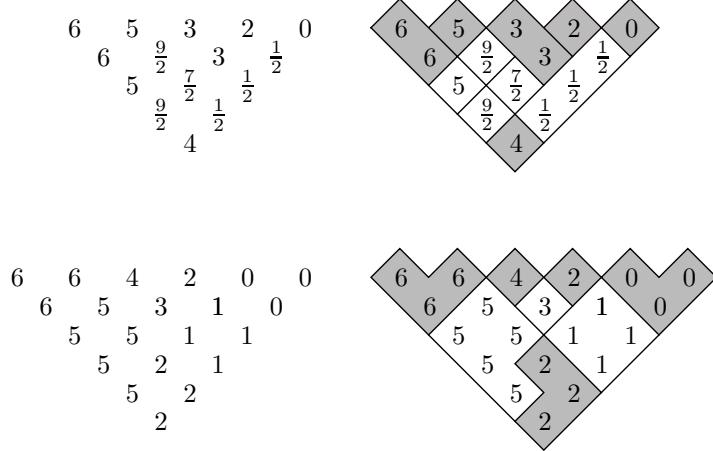


FIGURE 1. Tilings of GT-patterns.

Theorem 2.4 ([3, Theorem 1.5]). *Suppose that \mathcal{T} is the tiling of a GT-pattern \mathbf{x} . Then the dimension of the kernel of $A_{\mathcal{T}}$ is equal to the dimension of the minimal (dimensional) face of the GT-polytope containing \mathbf{x} .*

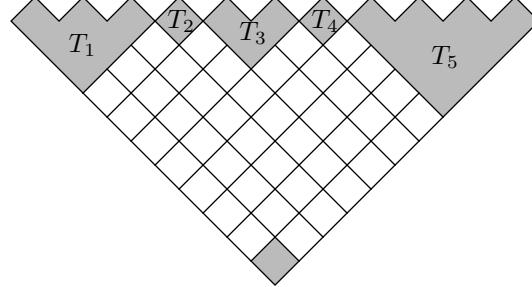
Theorem 2.4 was used in [3] to study the properties of vertices of GT-polytopes, establishing in particular that they can have arbitrarily large denominators. In the present note, we move to the opposite end of the face lattice and apply the tiling machinery to points in the interior of GT-polytopes. We use the notation $\text{int } P$ to denote the *relative interior* of a polytope P —that is, the interior of P with respect to the affine space that it spans. As an immediate corollary to Theorem 2.4, we get:

Corollary 2.5. *If $\mathbf{x} \in \text{int}(GT_{\lambda\beta})$, then $\dim GT_{\lambda\beta} = \dim \ker A_{\mathbf{x}}$.*

If $GT_{\lambda} \neq \emptyset$, the tilings of GT-patterns in the relative interior of GT_{λ} have an easy characterization. Observe that if a block of entries $x_{kr}, x_{k+1,r}, \dots, x_{\ell r}$ in the top row of a GT-pattern all have the same value κ_p , then the GT inequalities (3) require that the entries x_{ij} with $k \leq i \leq \ell - (r - j)$ all assume that same value κ_p . These are the entries that lie within the triangular region whose horizontal edge consists of the entries $x_{kr}, x_{k+1,r}, \dots, x_{\ell r}$ and whose diagonal edges run parallel to the diagonals of the GT-pattern. Let T_p be the tile of entries in this region. That is, if $\lambda = (\kappa_1^{v_1}, \dots, \kappa_m^{v_m}) \in \mathbb{R}^r$, put

$$T_p = \{(i, j) \in \mathcal{I} : \sum_{q=1}^{p-1} v_q < i \leq \sum_{q=1}^p v_q - (r - j)\}, \quad 1 \leq p \leq m.$$

Define the *generic interior tiling* \mathcal{T}_{λ} associated with λ to be the tiling consisting of these T_p 's together with a distinct tile for each entry not contained in one of the T_p 's. See Figure 2 for an example of a generic interior tiling when $(v_1, \dots, v_5) = (3, 1, 2, 1, 4)$. The shaded tile at the bottom of the pattern, and each of the unshaded free tiles, contains only a single entry.

FIGURE 2. The tiling of a point in the interior of GT_λ .

The next theorem characterizes when a GT-pattern \mathbf{x} has the generic interior tiling and gives the dimension of the GT-polytope containing \mathbf{x} . This is the main result from which Theorem 1.1 will follow.

Theorem 2.6. *Let $\lambda = (\kappa_1^{v_1}, \dots, \kappa_m^{v_m}) \in \mathbb{R}^r$.*

- (1) *A GT-pattern $\mathbf{x} \in GT_\lambda$ is in the relative interior of GT_λ if and only if the tiling of \mathbf{x} is the generic interior tiling \mathcal{T}_λ .*
- (2) *If \mathbf{x} has the generic interior tiling \mathcal{T}_λ and $m \geq 2$, then*

$$\dim GT(\mathbf{x}) = \binom{r-1}{2} - \sum_{p=1}^m \binom{v_p}{2},$$

(where we evaluate $\binom{1}{2} = 0$).

Proof. To prove part (1), note that the GT inequalities (3) include the facet-defining inequalities of GT_λ in X_r . A point \mathbf{x} lies in the relative interior of GT_λ if and only if \mathbf{x} satisfies with equality *only* those GT-inequalities that are satisfied with equality by every point in GT_λ .

We claim that the GT-inequalities satisfied with equality by every point in GT_λ are precisely the ones implied by forcing each entry in T_p to be equal to κ_p for $1 \leq p \leq m$. Since every point in GT_λ must satisfy those equalities, it remains only to exhibit a point $\mathbf{x} \in GT_\lambda$ satisfying *only* those equalities. That is, we need to construct an $\mathbf{x} \in GT_\lambda$ with the generic interior tiling.

For $1 \leq p \leq m-1$, consider the upper-left-to-lower-right diagonals that abut the tile T_p . Let S_p be the set of entries in these diagonals. Apply a total order to S_p by reading each diagonal from left to right, and then reading the diagonals themselves from left to right. Finally, construct \mathbf{x} by assigning values to the entries in S_p according to a strictly monotonically decreasing function mapping S_p into the open interval (κ_{p+1}, κ_p) . For example, in the portion depicted in Figure 3, we fill the tiles with strictly decreasing values in the order indicated by the arrows. It is easy to see that \mathbf{x} is a GT-pattern and that its tiling is \mathcal{T}_λ , so part (1) is proved.

To prove part (2), suppose that \mathbf{x} has the generic interior tiling. Then, by part (1), we have that $\mathbf{x} \in \text{int}(GT(\mathbf{x}))$. Thus, we can apply Corollary 2.5 to compute $\dim GT(\mathbf{x})$. The hypothesis that $m \geq 2$ implies that every free row of \mathbf{x} contains a free tile. Hence, each of the $r-2$ rows of the tiling matrix $A_{\mathbf{x}}$ contains a nonzero entry. Moreover, since every free tile of the generic interior tiling \mathcal{T}_λ contains only a single entry, every column of $A_{\mathbf{x}}$ contains only a single 1. Thus, $A_{\mathbf{x}}$ is in reduced row

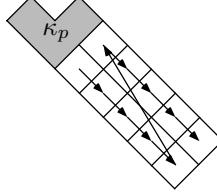


FIGURE 3. Filling of a portion of the generic interior tiling.

echelon form (perhaps after a suitable permutation of its columns, which amounts to re-indexing the free tiles). This means that

$$\begin{aligned} \dim GT(\mathbf{x}) &= \dim \ker A_{\mathbf{x}} \\ &= (\# \text{ of columns of } A_{\mathbf{x}}) - (\text{dimension of row span of } A_{\mathbf{x}}) \\ &= (\# \text{ of free tiles in } \mathcal{T}_{\lambda}) - (r - 2). \end{aligned}$$

Thus the computation reduces to finding the number of free tiles in \mathcal{T}_{λ} , which is easily done:

$$\binom{r}{2} - \sum_{p=1}^m (|T_p| - v_p) - 1 = \binom{r}{2} - \sum_{p=1}^m \binom{v_p}{2} - 1.$$

Hence,

$$\dim GT(\mathbf{x}) = \binom{r}{2} - \sum_{p=1}^m \binom{v_p}{2} - (r - 1) = \binom{r-1}{2} - \sum_{p=1}^m \binom{v_p}{2},$$

as claimed. \square

3. THE DEGREE OF STRETCHED KOSTKA COEFFICIENTS

In the previous section, we showed that if the interior points of $GT_{\lambda\beta}$ have the generic interior tiling, then the dimension of $GT_{\lambda\beta}$ is given by the expression in Theorem 1.1. To complete our proof of Theorem 1.1, it remains only to show that, if $\beta \lhd \lambda$, then the interior points of $GT_{\lambda\beta}$ have the generic interior tiling. To this end, we call upon a well-known fact from the theory of convex polytopes.

Lemma 3.1. *Given a polytope $P \subset \mathbb{R}^m$ and a linear map $\pi: \mathbb{R}^m \rightarrow \mathbb{R}^n$, we have that $\text{int}(\pi(P)) \subseteq \pi(\text{int } P)$.*

The converse containment is also true and easy to prove, but it is unnecessary for our purposes. It is also worth mentioning that this is the point at which the geometry of *convex* polyhedra is crucial to the argument. Lemma 3.1 does not hold for every polyhedral set, even if it is connected and full-dimensional. For completeness, we now give a proof of Lemma 3.1.

Proof of Lemma 3.1. Without loss of generality, we assume that $\dim P = m$ and $\dim \pi(P) = n$. Suppose that $y \notin \pi(\text{int } P)$. We show that $y \notin \text{int}(\pi(P))$ by exhibiting an “exit vector” \bar{y} such that, for every $\varepsilon > 0$, $y + \varepsilon\bar{y} \notin \pi(P)$. Since we assume that $\pi(P)$ is full-dimensional, this will prove the claim.

Since $y \notin \pi(\text{int } P)$, we have that $\pi^{-1}(y) \cap \text{int } P = \emptyset$. The sets $\pi^{-1}(y)$ and $\text{int } P$ are both convex, so there is a hyperplane H separating them. Let $\bar{x} \in \mathbb{R}^m$ be a

normal to H pointing away from P . Then, for every $x' \in \pi^{-1}(y)$ and $\varepsilon > 0$, we have that $x' + \varepsilon \bar{x} \notin P$. Let $\bar{y} = \pi(\bar{x})$. Suppose that $\varepsilon > 0$ and that $x \in \mathbb{R}^m$ is such that $\pi(x) = y + \varepsilon \bar{y}$. Then $x - \varepsilon \bar{x} \in \pi^{-1}(y)$, so $x = x - \varepsilon \bar{x} + \varepsilon \bar{x} \notin P$. In other words, $y + \varepsilon \bar{y} \notin \pi(P)$, proving the claim. \square

Let $D(\lambda) \subset \mathbb{R}^r$ be the image of GT_λ under the map $\mathbf{wt}: X_r \rightarrow \mathbb{R}^r$. Note that for $\beta \in D(\lambda)$, we have $GT_{\lambda\beta} = \mathbf{wt}^{-1}(\beta) \cap GT_\lambda$. It is well known that $GT_{\lambda\beta} \neq \emptyset$ if and only if $\beta \trianglelefteq \lambda$ [10]. For, if $\beta \trianglelefteq \lambda$, then $K_{\lambda\beta} > 0$ (see, e.g., [5, Exercise A.11]), so $GT_{\lambda\beta}$ contains an integral point. Conversely, if $GT_{\lambda\beta} \neq \emptyset$, then some integral multiple $nGT_{\lambda\beta} = GT_{n\lambda, n\beta}$ contains an integral point. Hence, $K_{n\lambda, n\beta} > 0$, so $n\beta \trianglelefteq n\lambda$. Since the relative order of two weights is not changed by scaling both by n , it follows that $\beta \trianglelefteq \lambda$. This establishes the following lemma:

Lemma 3.2. *Suppose that $\lambda \in \mathbb{Z}^r$ is a partition. Then $D(\lambda) = \{\beta \in \mathbb{R}^r : \beta \trianglelefteq \lambda\}$. Consequently, if $\beta \triangleleft \lambda$, then $\beta \in \text{int}(D(\lambda))$.*

Putting together the preceding results, we are now ready to prove Theorem 1.1 from page 2.

Proof of Theorem 1.1. The claim is trivial if $r = 1$, so suppose that $r \geq 2$. From Lemmas 3.2 and 3.1, we have that

$$\beta \in \text{int}(D(\lambda)) \subset \mathbf{wt}(\text{int}(GT_\lambda)).$$

Hence,

$$GT_{\lambda\beta} \cap \text{int}(GT_\lambda) = \mathbf{wt}^{-1}(\beta) \cap \text{int}(GT_\lambda) \neq \emptyset.$$

Choose $\mathbf{x} \in GT_{\lambda\beta} \cap \text{int}(GT_\lambda)$. By part (1) of Theorem 2.6, \mathbf{x} has the generic interior tiling \mathcal{T}_λ . Note that since $\beta \triangleleft \lambda$, we must have that $m \geq 2$, for otherwise the GT-inequalities would force $\beta_1 = \lambda_1$. Hence, we may apply part (2) of Theorem 2.6 to compute the dimension of $GT(\mathbf{x}) = GT_{\lambda\beta}$. Since the dimension of $GT_{\lambda\beta}$ is the degree of $\mathcal{K}_{\lambda\beta}(n)$ [13, Theorem 4.6.25], this yields

$$\begin{aligned} \deg \mathcal{K}_{\lambda\beta}(n) &= \dim GT_{\lambda\beta} = \dim GT(\mathbf{x}) \\ &= \binom{r-1}{2} - \sum_{p=1}^m \binom{v_p}{2}, \end{aligned}$$

as claimed. \square

ACKNOWLEDGEMENTS

The author wishes to thank Christophe Tollu and Frederic Toumazet for exposing him to their conjecture, and Jesús De Loera for helpful comments and suggestions. The author also thanks the anonymous referees for suggestions improving the exposition.

REFERENCES

- [1] A. D. Berenstein and A. V. Zelevinsky, *When is the multiplicity of a weight equal to 1?*, Funktsional. Anal. i Prilozhen. **24** (1990), no. 4, 1–13, 96.
- [2] S. Billey, V. Guillemin, and E. Rassart, *A vector partition function for the multiplicities of $\mathfrak{sl}_k\mathbb{C}$* , J. Algebra **278** (2004), no. 1, 251–293.
- [3] J. A. De Loera and T. B. McAllister, *Vertices of Gelfand-Tsetlin polytopes*, Discrete Comput. Geom. **32** (2004), no. 4, 459–470, arXiv:math.CO/0309329.
- [4] W. Fulton, *Young tableaux*, London Mathematical Society Student Texts, vol. 35, Cambridge University Press, Cambridge, 1997, With applications to representation theory and geometry.

- [5] W. Fulton and J. Harris, *Representation theory*, Graduate Texts in Mathematics, vol. 129, Springer-Verlag, New York, 1991, A first course, Readings in Mathematics.
- [6] I. M. Gelfand and M. L. Tsetlin, *Finite-dimensional representations of the group of unimodular matrices*, Doklady Akad. Nauk SSSR (N.S.) **71** (1950), 825–828, English translation in *Collected Papers, Vol. II*, Springer, Berlin, 1988, 653–656.
- [7] R. C. King, C. Tollu, and F. Toumazet, *Stretched Littlewood-Richardson and Kostka coefficients*, Symmetry in physics, CRM Proc. Lecture Notes, vol. 34, Amer. Math. Soc., Providence, RI, 2004, pp. 99–112.
- [8] ———, *The hive model and the factorisation of Kostka coefficients*, Séminaire Lotharingien de Combinatoire **54A** (2006), Article B54Ah.
- [9] A. N. Kirillov, *Ubiquity of Kostka polynomials*, Physics and combinatorics 1999 (Nagoya), World Sci. Publishing, River Edge, NJ, 2001, arXiv:math.QA/9912094, pp. 85–200.
- [10] A. N. Kirillov and A. D. Berenstein, *Groups generated by involutions, Gel'fand-Tsetlin patterns, and combinatorics of Young tableaux*, Algebra i Analiz **7** (1995), no. 1, 92–152, translation in *St. Petersburg Math. J.* **7** (1996), no. 1, 77–127.
- [11] A. N. Kirillov and N. Y. Reshetikhin, *The Bethe ansatz and the combinatorics of Young tableaux*, Zap. Nauchn. Sem. Leningrad. Otdel. Mat. Inst. Steklov. (LOMI) **155** (1986), no. Differentsialnaya Geometriya, Gruppy Li i Mekh. VIII, 65–115, 194, translation in *J. Soviet Math.* **41** (1988), no. 2, 925–955.
- [12] R. P. Stanley, *Some combinatorial aspects of the Schubert calculus*, Combinatoire et représentation du groupe symétrique (Actes Table Ronde CNRS, Univ. Louis-Pasteur Strasbourg, Strasbourg, 1976), Springer, Berlin, 1977, pp. 217–251. Lecture Notes in Math., Vol. 579.
- [13] ———, *Enumerative combinatorics. Vol. 1*, Cambridge Studies in Advanced Mathematics, vol. 49, Cambridge University Press, Cambridge, 1997, With a foreword by Gian-Carlo Rota, Corrected reprint of the 1986 original.
- [14] ———, *Enumerative combinatorics. Vol. 2*, Cambridge Studies in Advanced Mathematics, vol. 62, Cambridge University Press, Cambridge, 1999, With a foreword by Gian-Carlo Rota and appendix 1 by Sergey Fomin.

DEPARTMENT OF MATHEMATICS AND COMPUTER SCIENCE, EINDHOVEN UNIVERSITY OF TECHNOLOGY, P.O. BOX 513, 5600 MB EINDHOVEN, THE NETHERLANDS

E-mail address: `tmcallis@win.tue.nl`